

REMARKS

Claims 1-52 remain pending in the application. Reconsideration is respectfully requested in light of the following remarks.

Section 102(e) Rejection:

The Office Action rejected claims 1-52 under 35 U.S.C. § 102(e) as being anticipated by Strait (U.S. Publication 2005/0015437). Applicant traverses this rejection for at least the following reason.

In regard to claim 1, contrary to the Examiner's assertion, Strait does not anticipate a master node configured to manage a grid comprising one or more compute nodes; and a node configured to send the master node information about compute node configuration of the node in accordance with one or more peer-to-peer platform protocols.

The Examiner cites paragraphs [0018]-[0019] and [0022]-[0023] of Strait. Strait is directed at “peer to peer job monitoring and control in grid computing systems” (Strait, Title). Paragraph [0018] describes FIG. 1, which is a “typical system for distributing a workload from a plurality of submitters to a plurality of clients for processing.” Strait’s FIG. 1 is actually exemplary of the general architecture of a prior art grid system. Figures 1 and 2 of the instant application and the accompanying description thereof at page 2, lines 6-30 illustrate and describe exemplary prior art cluster grids and grid farms, in which a management node and “master node” are more or less analogous to Strait’s “centralized server” of FIG. 1, access nodes and job submitter node are more or less analogous to Strait’s “submitting clients” of FIG. 1, and compute nodes are more or less analogous to Strait’s “processing clients” of FIG. 1. However, nowhere in paragraph [0018] does Strait disclose anything like a node configured to send the master node information about compute node configuration of the node in accordance with one or more peer-to-peer platform protocols. The only things that Strait indicates may be sent

from a client to the centralized server over the “communication links” are requests for processing services from the submitting clients, and indirectly notification of completion of processing and final results from the processing clients.

Paragraph [0019] of Strait describes FIG. 2, which illustrates “additional communication paths opened by the subject invention for communication between clients.” Strait discloses (emphasis added): “these communication paths are established by the clients, with no participation by the centralized server 1 shown in FIG. 1.” In paragraph [0006], Strait actually teaches (emphasis added): “In the disclosed invention, a centralized server as previously used for dispatch and management of batch jobs **remains unmodified, as in prior art.**” In paragraph [0018], referring to FIG. 1, Strait discloses “This communications structure remains unchanged for the present invention, and continues to fulfill these roles, and centralized server 1 continues to fill the important role of workload balancing among the clients.” In other words, the “centralized server” shown in and described for Strait’s FIG. 2 is the same as the prior art centralized server shown in and described for Strait’s FIG. 1. Thus, Strait’s disclosed system as illustrated in and described for FIG. 2 does not and cannot add anything to the prior art system of FIG. 1 that would teach a node configured to send the master node information about compute node configuration of the node in accordance with one or more peer-to-peer platform protocols. Strait only discloses that the centralized server “continues to fill the role of workload balancing among the clients.”

Regarding paragraphs [0022]-[0023], paragraph [0022] mentions that a [submitting] client 2 submits a request to the “batch server 1”, i.e. Strait’s centralized server, which again is not changed from the prior art centralized server. Strait teaches “This request may be in the form of a control file, such as shown in FIG. 4. This control file is typical of prior art systems. This example control file specifies a 2 step job requiring 2 processing clients, but any number of job steps may be specified.” The paragraph clearly does not teach a node configured to send the master node information about compute node configuration of the node in accordance with one or more peer-to-peer platform protocols. Furthermore, Strait’s request/control file, which “specifies a

[multi]step job requiring [multiple] processing clients”, is clearly not analogous to compute node configuration of a node, as is recited in Applicant’s claim 1.

Paragraph [0023] of Strait discloses “To accomplish the monitoring and control disclosed by this invention, the request by client 2 to batch server 1 must be modified to cause the processing clients to start a monitoring process ahead of the actual batch job step. This is accomplished by modifying the control file of FIG. 4 as shown in FIG. 5.” At paragraph [0029], Strait discloses (emphasis added): “The required modifications shown in FIG. 5 will be made on client 2 prior to submission of the request to the centralized server 1.” Again, Strait’s request/control file, which is modified on the submitting node prior to submission to Strait’s centralized server, is clearly not analogous to compute node configuration of a node, as is recited in Applicant’s claim 1.

In further regard to claim 1, contrary to the Examiner’s assertion, Strait does not anticipate wherein the master node is configured to determine from the information about compute node configuration that the compute node configuration of the node needs to be updated.

The Examiner again cites paragraphs [0018] and [0022]-[0023] of Strait. Applicant has discussed these paragraphs above. Nothing in these paragraphs teaches or suggests anything like a master node configured to determine from the information about compute node configuration that the compute node configuration of the node needs to be updated. Nothing in these paragraphs even teaches a master node receiving information about compute node configuration from a node. Paragraphs [0018] and [0022]-[0023] only teach that the centralized server receives a request/control file, which “specifies a [multi]step job requiring [multiple] processing clients”, from a submitting client, and paragraph [0018] only suggests that the centralized server may receive results or notifications of completions from processing clients. Furthermore, as noted, Strait clearly teaches that the “centralized server” remains unchanged from the prior art, and nothing in the prior art as disclosed by Strait, or elsewhere in Strait’s teachings, teaches or suggests anything like the centralized server determining from information about compute node

configuration received from a node that a compute node configuration of the node needs to be updated.

In further regard to claim 1, contrary to the Examiner's assertion, Strait does not anticipate wherein the master node is configured to send update information for the compute node configuration to the node in accordance with the one or more peer-to-peer platform protocols.

The Examiner cites paragraphs [0027]-[0030] of Strait. Paragraph [0027] simply mentions an “optional secret security key used by processing clients to authenticate their access to the submitter system” which may be added in an argument to the control file. At paragraph [0029], Strait discloses that the required modifications to the control file are made on the submitting client 2 prior to submission of the request to the centralized server. Strait’s request/control file, which is modified on the submitting node prior to submission to Strait’s centralized server, optionally with the “secret security key” added, is clearly not analogous to update information for the compute node configuration, as is recited in Applicant’s claim 1.

In addition, Strait’s centralized server simply receives the modified control file from a submitting node and distributes the “steps” indicated therein to multiple processing clients. Again, Strait clearly teaches that the centralized server “remains unchanged” from a prior art centralized server, and thus performs the task of distributing steps to processing nodes without any modifications in Strait’s system. Strait does not teach, as noted above, that the centralized server determines, from the control file or from the modifications to the control file described in paragraphs [0022]-[0028], anything about compute node configuration, e.g. that the compute node configuration of a node needs to be updated, and sends update information for the compute node configuration to the node that sent the master node information about compute node configuration of the node.

Paragraph [0030] of Strait simply teaches, to begin the job submission process, “the submitter, on client 2, invokes a program that will prepare to accept communications from the processing clients...and that will submit the modified job requests to centralized server 1 for processing.” The paragraph does not teach anything like the master node sending update information for the compute node configuration to the node.

In further regard to claim 1, Strait teaches a “centralized server as previously used for dispatch and management of batch jobs” that remains unmodified, as in prior art, and thus does not anticipate Applicant’s claim 1.

Strait discloses (emphasis added) “A solution for improved monitoring and control of jobs in grid and batch computing systems provides a centralized server's batch manager which is only responsible for workload balancing and job initiation and completion, all other command and status information are communicated directly between the plurality of submitter's systems and the plurality of client systems that are processing their respective workloads.” (Strait, Abstract). In paragraph [0004], Strait teaches a prior art “centralized server”:

A typical implementation of a parallel processing system allows multiple clients to each submit multiple job steps to be distributed among a pool of clients for processing. This is typically accomplished by having a centralized server receive all such requests, and prioritize and distribute them to a pool of client systems for processing. The centralized server is responsible for workload balancing among the clients, and for the commands to the client systems necessary to start and maintain jobs, and for monitoring activity and notifying the submitter of status, such as completion of the job steps on each client.

In paragraph [0006] Strait teaches (emphasis added): “In the disclosed invention, a centralized server as previously used for dispatch and management of batch jobs **remains unmodified, as in prior art.**”

Figures 1 and 2 of the instant application and the accompanying description thereof at page 2, lines 6-30 illustrate and describe exemplary prior art cluster grids and grid farms, in which a “management node” and “master node” are more or less analogous

to Strait's "centralized server." Applicant's claim 1 clearly recites additional actions performed by a master node that are not described for a prior art centralized server in Strait or for a prior art master node/management node in Applicant's background section. For example, claim 1 recites a master node configured to determine from the information [received from the node] about compute node configuration that the compute node configuration of the node needs to be updated, and further configured to send update information for the compute node configuration to the node in accordance with the one or more peer-to-peer platform protocols. **Since Strait clearly teaches "a centralized server as previously used for dispatch and management of batch jobs remains unmodified, as in prior art", Strait's system does not include the limitations as recited in Applicant's claim 1, and Strait's teachings does not anticipate Applicant's claim 1.**

In further regard to claim 1, Strait does not anticipate Applicant's claim 1 when viewed as a whole.

Strait discloses, in paragraph [0006], that a centralized server as previously used for dispatch and management of batch jobs remains unmodified, as in prior art. However, Strait discloses a system that provides for improved communication and feedback between the originator of jobs and the client systems processing those jobs, such that the submitter is able to receive real time feedback on job progress, and to send real-time commands to alter the continued operation of said jobs. The submitter application of batch jobs applications allows interacting with batch jobs during their execution. Software layers are added at each submitter's system and at each batch client where a job step is processed. These additional software layers permit communication directly between the submitter's client system and each processing client, without use of a communications channel through the centralized batch manager.

In contrast, Applicant's claim 1 recites a grid computing system, comprising a master node configured to manage a grid comprising one or more compute nodes and a node configured to send the master node information about compute node configuration

of the node in accordance with one or more peer-to-peer platform protocols. The master node is configured to determine from the information about compute node configuration that the compute node configuration of the node needs to be updated, and send update information for the compute node configuration to the node in accordance with the one or more peer-to-peer platform protocols. **Nowhere in Strait is this subject matter taught as it is recited in Applicant's claim 1.**

Applicant reminds the Examiner that anticipation requires the presence in a single prior art reference disclosure of each and every element of the claimed invention, arranged as in the claim. M.P.E.P 2131; *Lindemann Maschinenfabrik GmbH v. American Hoist & Derrick Co.*, 221 USPQ 481, 485 (Fed. Cir. 1984). The identical invention must be shown in as complete detail as is contained in the claims. *Richardson v. Suzuki Motor Co.*, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). Nowhere does the Strait reference disclose “each and every element of the claimed invention” (claim 1 of the instant application) as arranged in the claim. For example, Strait does not disclose a node configured to send the master node information about compute node configuration of the node. As another example, Strait does not disclose a master node configured to determine from the information about compute node configuration that the compute node configuration of the node needs to be updated. As another example, Strait does not disclose a master node configured to send update information for the compute node configuration to the node. Furthermore, even if Strait did disclose one or more of the above elements, nowhere does Strait disclose the above elements arranged as in claim 1. For at least the reasons given above, Strait clearly does not anticipate Applicant's claim 1.

Thus, for at least the reasons presented above, the rejection of claim 1 is not supported by the cited art and removal thereof is respectfully requested.

In regard to claim 21, the Examiner has not provided a proper *prima facie* rejection of the claim. The Examiner rejected claim 21 with claim 1. However, claims 1 and 21 are of different scope. For example, claim 1 recites a master node configured to manage a grid comprising one or more compute nodes. Claim 21 recites no such

limitation. As another example, claim 21 recites a system configured to participate as a compute node in a grid configured to, if the compute node configuration of the system is not up-to-date: obtain update information for the compute node configuration from the node in accordance with the one or more peer-to-peer platform protocols; and update the compute node configuration of the system in accordance with the update information. Claim 1 does not recite this subject matter as recited in claim 21.

However, in further regard to claim 21, Strait does not anticipate a system configured to participate as a compute node in a grid comprising one or more compute nodes, comprising: a processor; and a memory comprising program instructions, wherein the program instructions are executable by the processor to: communicate with a node on a network in accordance with one or more peer-to-peer platform protocols to determine if compute node configuration of the system is up-to-date.

In the rejection of claim 1, the Examiner cites paragraphs [0018]-[0019] and [0022]-[0023] of Strait in regard to similar (but not identical) subject matter. Furthermore, in the rejection of claim 36, the Examiner cites paragraphs [0018] and [0022]-[0023] of Strait in regard to similar (but not identical) subject matter. Strait is directed at “peer to peer job monitoring and control in grid computing systems” (Strait, Title). Paragraph [0018] describes FIG. 1, which is a “typical system for distributing a workload from a plurality of submitters to a plurality of clients for processing.” Strait’s FIG. 1 is actually exemplary of the general architecture of a prior art grid system. However, nowhere in paragraph [0018] does Strait disclose anything like a system configured to participate as a compute node in a grid comprising one or more compute nodes configured to communicate with a node on a network to determine if compute node configuration of the system is up-to-date. The only things that paragraph [0018] indicates may be sent from a client to another node over the “communication links” are requests for processing services from the submitting clients to the centralized server, and notification of completion of processing and final results from the central server to the submitting clients.

Paragraph [0019] of Strait describes FIG. 2, which illustrates “additional communication paths opened by the subject invention for communication between clients.” However, nowhere in paragraph [0019] does Strait disclose anything like a system configured as a compute node in a grid configured to communicate with a node on a network to determine if compute node configuration of the system is up-to-date. The only things that paragraph [0019] indicates may be sent from a client to another node over the “communication links” are requests for processing services from the submitting clients to the centralized server, and assignments for the requests from the centralized server to the processing clients.

Regarding paragraphs [0022]-[0023], paragraph [0022] mentions that a [submitting] client 2 submits a request to the “batch server 1”, i.e. Strait’s centralized server. Strait teaches “This request may be in the form of a control file, such as shown in FIG. 4. This control file is typical of prior art systems. This example control file specifies a 2 step job requiring 2 processing clients, but any number of job steps may be specified.” The paragraph clearly does not teach a system configured as a compute node in a grid configured to communicate with a node on a network to determine if compute node configuration of the system is up-to-date. Furthermore, Strait’s request/control file, which “specifies a [multi]step job requiring [multiple] processing clients”, is clearly not analogous to communications from a compute node to another node on a network to determine if compute node configuration of the system is up-to-date, as is recited in Applicant’s claim 21.

Paragraph [0023] of Strait discloses “To accomplish the monitoring and control disclosed by this invention, the request by client 2 to batch server 1 must be modified to cause the processing clients to start a monitoring process ahead of the actual batch job step. This is accomplished by modifying the control file of FIG. 4 as shown in FIG. 5.” At paragraph [0029], Strait discloses (emphasis added): “The required modifications shown in FIG. 5 will be made on client 2 prior to submission of the request to the centralized server 1.” The paragraph clearly does not teach a system configured as a

compute node in a grid configured to communicate with a node on a network to determine if compute node configuration of the system is up-to-date. Again, Strait's request/control file, which is modified on the submitting node prior to submission to Strait's centralized server, is clearly not analogous to communications from a compute node to another node on a network to determine if compute node configuration of the system is up-to-date, as is recited in Applicant's claim 21.

In further regard to claim 21, Strait does not anticipate a system configured to participate as a compute node in a grid configured to, if the compute node configuration of the system is not up-to-date: obtain update information for the compute node configuration from the node in accordance with the one or more peer-to-peer platform protocols.

In the rejection of claim 1, the Examiner cites paragraphs [0018] and [0022]-[0023] of Strait in regard to similar (but not identical) subject matter. Furthermore, in the rejection of claim 36, the Examiner cites paragraphs [0018] and [0022]-[0023] of Strait in regard to similar (but not identical) subject matter.

Applicant has discussed these paragraphs above. Nothing in these paragraphs teaches or suggests anything like a system configured to participate as a compute node in a grid configured to, if the compute node configuration of the system is not up-to-date: obtain update information for the compute node configuration from the node. Paragraphs [0018] and [0022]-[0023] only teach that the centralized server receives a request/control file, which "specifies a [multi]step job requiring [multiple] processing clients", from a submitting client, and paragraph [0018], directed at prior art systems, only suggests that the centralized server may receive results or notifications of completions from processing clients.

Regarding paragraphs [0022]-[0023], paragraph [0022] mentions that a [submitting] client 2 submits a request to the "batch server 1", i.e. Strait's centralized server. Strait teaches "This request may be in the form of a control file, such as shown in

FIG. 4. This control file is typical of prior art systems. This example control file specifies a 2 step job requiring 2 processing clients, but any number of job steps may be specified.” The paragraph clearly does not teach a system configured to participate as a compute node in a grid configured to, if the compute node configuration of the system is not up-to-date: obtain update information for the compute node configuration from the node. Furthermore, Strait’s request/control file, which “specifies a [multi]step job requiring [multiple] processing clients”, is clearly not analogous to update information for the compute node configuration, as is recited in Applicant’s claim 21.

Paragraph [0023] of Strait discloses “To accomplish the monitoring and control disclosed by this invention, the request by client 2 to batch server 1 must be modified to cause the processing clients to start a monitoring process ahead of the actual batch job step. This is accomplished by modifying the control file of FIG. 4 as shown in FIG. 5.” At paragraph [0029], Strait discloses (emphasis added): “The required modifications shown in FIG. 5 will be made on client 2 prior to submission of the request to the centralized server 1.” The paragraph clearly does not teach a system configured to participate as a compute node in a grid configured to, if the compute node configuration of the system is not up-to-date: obtain update information for the compute node configuration from the node. Again, Strait’s request/control file, which is modified on the submitting node prior to submission to Strait’s centralized server, is clearly not analogous update information for the compute node configuration, as is recited in Applicant’s claim 21.

In further regard to claim 21, Strait does not anticipate a system configured to participate as a compute node in a grid configured to, if the compute node configuration of the system is not up-to-date: obtain update information for the compute node configuration from the node in accordance with the one or more peer-to-peer platform protocols and update the compute node configuration of the system in accordance with the update information.

Claim 1 does not recite the limitation update the compute node configuration of the system in accordance with the update information. However, in the rejection of claim 36, the Examiner cites paragraphs [0018] and [0022]-[0023] of Strait in regard to similar (but not identical) subject matter. Applicant has discussed these paragraphs above. Nothing in these paragraphs teaches or suggests anything like the above limitations. Strait does not teach in these citations a system configured as a compute node configured to obtain update information for the compute node configuration from another node and update the compute node configuration of the system in accordance with the update information.

In further regard to claim 21, Strait does not anticipate Applicant's claim 21 when viewed as a whole.

Strait discloses a system that provides for improved communication and feedback between the originator of jobs and the client systems processing those jobs, such that the submitter is able to receive real time feedback on job progress, and to send real-time commands to alter the continued operation of said jobs. The submitter application of batch jobs applications allows interacting with batch jobs during their execution. Software layers are added at each submitter's system and at each batch client where a job step is processed. These additional software layers permit communication directly between the submitter's client system and each processing client, without use of a communications channel through the centralized batch manager.

In contrast, Applicant's claim 21 recites a system configured to participate as a compute node in a grid comprising one or more compute nodes, comprising: a processor and a memory comprising program instructions, wherein the program instructions are executable by the processor to communicate with a node on a network in accordance with one or more peer-to-peer platform protocols to determine if compute node configuration of the system is up-to-date; and if the compute node configuration of the system is not up-to-date, obtain update information for the compute node configuration from the node in accordance with the one or more peer-to-peer platform protocols and update the

compute node configuration of the system in accordance with the update information. **Nowhere in Strait is this subject matter taught as it is recited in Applicant's claim 21.**

Applicant reminds the Examiner that anticipation requires the presence in a single prior art reference disclosure of each and every element of the claimed invention, arranged as in the claim. M.P.E.P 2131; *Lindemann Maschinenfabrik GmbH v. American Hoist & Derrick Co.*, 221 USPQ 481, 485 (Fed. Cir. 1984). The identical invention must be shown in as complete detail as is contained in the claims. *Richardson v. Suzuki Motor Co.*, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). Nowhere does the Strait reference disclose “each and every element of the claimed invention” (claim 21 of the instant application) as arranged in the claim. For example, Strait does not disclose a system configured to participate as a compute node on a grid and configured to communicate with a node on a network to determine if compute node configuration of the system is up-to-date. As another example, Strait does not disclose if the compute node configuration of the system is not up-to-date, the system is configured to obtain update information for the compute node configuration from the node and update the compute node configuration of the system in accordance with the update information. Furthermore, even if Strait did disclose one or more of the above elements, nowhere does Strait disclose the above elements arranged as in claim 21. For at least the reasons given above, Strait clearly does not anticipate Applicant's claim 21.

Thus, for at least the reasons presented above, the rejection of claim 21 is not supported by the cited art and removal thereof is respectfully requested.

In regard to claim 29, the Examiner has not provided a proper *prima facie* rejection of the claim. The Examiner rejected claim 29 with claim 1. However, claims 1 and 29 are of different scope. For example, claim 1 recites a master node configured to manage a grid comprising one or more compute nodes. Claim 29 recites no such limitation. As another example, claim 29 recites a system comprising a processor and a memory comprising program instructions, wherein the program instructions are

executable by the processor to receive information about compute node configuration of a node configured to participate as a compute node in a grid in accordance with one or more peer-to-peer platform protocols. Claim 1 does not recite this specific limitation.

However, in further regard to claim 29, Strait does not anticipate a system comprising a processor and a memory comprising program instructions, wherein the program instructions are executable by the processor to receive information about compute node configuration of a node configured to participate as a compute node in a grid in accordance with one or more peer-to-peer platform protocols.

In the rejection of claim 1, regarding the subject matter *a node configured to send the master node information about compute node configuration of the node in accordance with one or more peer-to-peer platform protocols* as recited in claim 1, the Examiner cites paragraphs [0018]-[0019] and [0022]-[0023] of Strait in regard to similar (but not identical) subject matter. Furthermore, in the rejection of claim 36, the Examiner cites paragraphs [0018] and [0022]-[0023] of Strait in regard to similar (but not identical) subject matter. Strait is directed at “peer to peer job monitoring and control in grid computing systems” (Strait, Title). Paragraph [0018] describes FIG. 1, which is a “typical system for distributing a workload from a plurality of submitters to a plurality of clients for processing.” Strait’s FIG. 1 is actually exemplary of the general architecture of a prior art grid system. However, nowhere in paragraph [0018] does Strait disclose anything like a system comprising a processor and a memory comprising program instructions, wherein the program instructions are executable by the processor to receive information about compute node configuration of a node configured to participate as a compute node in a grid. The only things that paragraph [0018] indicates may be sent from a client to another node over the “communication links” are requests for processing services from the submitting clients to the centralized server, and notification of completion of processing and final results from the central server to the submitting clients.

Paragraph [0019] of Strait describes FIG. 2, which illustrates “additional communication paths opened by the subject invention for communication between clients.” However, nowhere in paragraph [0019] does Strait disclose anything like a system comprising a processor and a memory comprising program instructions, wherein the program instructions are executable by the processor to receive information about compute node configuration of a node configured to participate as a compute node in a grid. The only things that paragraph [0019] indicates may be sent from a client to another node over the “communication links” are requests for processing services from the submitting clients to the centralized server, and assignments for the requests from the centralized server to the processing clients.

Regarding paragraphs [0022]-[0023], paragraph [0022] mentions that a [submitting] client 2 submits a request to the “batch server 1”, i.e. Strait’s centralized server. Strait teaches “This request may be in the form of a control file, such as shown in FIG. 4. This control file is typical of prior art systems. This example control file specifies a 2 step job requiring 2 processing clients, but any number of job steps may be specified.” The paragraph clearly does not teach a system comprising a processor and a memory comprising program instructions, wherein the program instructions are executable by the processor to receive information about compute node configuration of a node configured to participate as a compute node in a grid. Furthermore, Strait’s request/control file, which “specifies a [multi]step job requiring [multiple] processing clients”, is clearly not analogous to information about compute node configuration of a node, as is recited in Applicant’s claim 29.

Paragraph [0023] of Strait discloses “To accomplish the monitoring and control disclosed by this invention, the request by client 2 to batch server 1 must be modified to cause the processing clients to start a monitoring process ahead of the actual batch job step. This is accomplished by modifying the control file of FIG. 4 as shown in FIG. 5.” At paragraph [0029], Strait discloses (emphasis added): “The required modifications shown in FIG. 5 will be made on client 2 prior to submission of the request to the centralized server 1.” The paragraph clearly does not teach a system comprising a

processor and a memory comprising program instructions, wherein the program instructions are executable by the processor to receive information about compute node configuration of a node configured to participate as a compute node in a grid. Again, Strait's request/control file, which is modified on the submitting node prior to submission to Strait's centralized server, is clearly not analogous to information about compute node configuration of a node, as is recited in Applicant's claim 29.

In further regard to claim 29, Strait does not anticipate a system comprising a processor and a memory comprising program instructions, wherein the program instructions are executable by the processor to determine from the information about compute node configuration that the compute node configuration of the node needs to be updated.

In the rejection of claim 1, the Examiner again cites paragraphs [0018] and [0022]-[0023] of Strait in regard to similar subject matter. Nothing in these paragraphs teaches or suggests anything like a system configured to determine from the information about compute node configuration that the compute node configuration of the node needs to be updated. Nothing in these paragraphs even teaches a system receiving information about compute node configuration from a node. Paragraphs [0018] and [0022]-[0023] only teach that the centralized server receives a request/control file, which "specifies a [multi]step job requiring [multiple] processing clients", from a submitting client, and paragraph [0018] only suggests that the centralized server may receive results or notifications of completions from processing clients. Nothing in the prior art, or in Strait's teachings, teaches or suggests anything like a system determining from information about compute node configuration received from a node that a compute node configuration of the node needs to be updated.

In further regard to claim 29, Strait does not anticipate a system comprising a processor and a memory comprising program instructions, wherein the program instructions are executable by the processor to send update information for the

compute node configuration to the node in accordance with the one or more peer-to-peer platform protocols.

In the rejection of claim 1, the Examiner cites paragraphs [0027]-[0030] of Strait in regard to similar subject matter. Paragraph [0027] simply mentions an “optional secret security key used by processing clients to authenticate their access to the submitter system” which may be added in an argument to the control file. At paragraph [0029], Strait discloses that the required modifications to the control file are made on the submitting client 2 prior to submission of the request to the centralized server. Strait’s request/control file, which is modified on the submitting node prior to submission to Strait’s centralized server, optionally with the “secret security key” added, is clearly not analogous to update information for the compute node configuration, as is recited in Applicant’s claim 29.

In addition, Strait’s centralized server simply receives the modified control file from a submitting node and distributes the “steps” indicated therein to multiple processing clients. Strait does not teach that the centralized server determines, from the control file or from the modifications to the control file described in paragraphs [0022]-[0028], anything about compute node configuration, e.g. that the compute node configuration of a node needs to be updated, and sends update information for the compute node configuration to the node that sent the system information about compute node configuration of the node.

Paragraph [0030] of Strait simply teaches, to begin the job submission process, “the submitter, on client 2, invokes a program that will prepare to accept communications from the processing clients...and that will submit the modified job requests to centralized server 1 for processing.” The paragraph does not teach anything like a system sending update information for compute node configuration to the node.

In further regard to claim 29, Strait does not anticipate Applicant’s claim 29 when viewed as a whole.

Strait discloses a system that provides for improved communication and feedback between the originator of jobs and the client systems processing those jobs, such that the submitter is able to receive real time feedback on job progress, and to send real-time commands to alter the continued operation of said jobs. The submitter application of batch jobs applications allows interacting with batch jobs during their execution. Software layers are added at each submitter's system and at each batch client where a job step is processed. These additional software layers permit communication directly between the submitter's client system and each processing client, without use of a communications channel through the centralized batch manager.

In contrast, Applicant's claim 29 recites a system, comprising a processor and a memory comprising program instructions, wherein the program instructions are executable by the processor to receive information about compute node configuration of a node configured to participate as a compute node in a grid in accordance with one or more peer-to-peer platform protocols, determine from the information about compute node configuration that the compute node configuration of the node needs to be updated, and send update information for the compute node configuration to the node in accordance with the one or more peer-to-peer platform protocols **Nowhere in Strait is this subject matter taught as it is recited in Applicant's claim 29.**

Applicant reminds the Examiner that anticipation requires the presence in a single prior art reference disclosure of each and every element of the claimed invention, arranged as in the claim. M.P.E.P 2131; *Lindemann Maschinenfabrik GmbH v. American Hoist & Derrick Co.*, 221 USPQ 481, 485 (Fed. Cir. 1984). The identical invention must be shown in as complete detail as is contained in the claims. *Richardson v. Suzuki Motor Co.*, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). Nowhere does the Strait reference disclose "each and every element of the claimed invention" (claim 21 of the instant application) as arranged in the claim. For example, Strait does not disclose a system configured to receive information about compute node configuration of a node configured to participate as a compute node in a grid. As another example, Strait does not disclose a system

configured to determine from the information about compute node configuration that the compute node configuration of the node needs to be updated and send update information for the compute node configuration to the node in accordance with the one or more peer-to-peer platform protocols. Furthermore, even if Strait did disclose one or more of the above elements, nowhere does Strait disclose the above elements arranged as in claim 29. For at least the reasons given above, Strait clearly does not anticipate Applicant's claim 29.

Thus, for at least the reasons presented above, the rejection of claim 29 is not supported by the cited art and removal thereof is respectfully requested.

In regard to claim 9, claim 9 is a method claim that recites similar subject matter to that recited in claim 1. Therefore, Applicant traverses the rejection of claim 9 for at least the reasons given above in regard to claim 1.

In regard to claim 15, claim 15 is a computer-accessible storage medium claim that recites similar subject matter to that recited in claim 1. Therefore, Applicant traverses the rejection of claim 15 for at least the reasons given above in regard to claim 1.

In regard to claim 36, claim 36 is a means plus function claim that recites similar subject matter to that recited in claim 21. Therefore, Applicant traverses the rejection of claim 36 for at least the reasons given above in regard to claim 21.

In regard to claim 37, claim 37 is a method claim that recites similar subject matter to that recited in claim 21. Therefore, Applicant traverses the rejection of claim 37 for at least the reasons given above in regard to claim 21.

In regard to claim 45, claim 45 is a computer-accessible storage medium claim that recites similar subject matter to that recited in claim 21. Therefore, Applicant

traverses the rejection of claim 45 for at least the reasons given above in regard to claim 21.

Applicant also asserts that the rejection of numerous ones of the dependent claims is further unsupported by the cited art. However, since the rejection has been shown to be unsupported for the independent claims, a further discussion of the dependent claims is not necessary at this time.

CONCLUSION

Applicant submits the application is in condition for allowance, and notice to that effect is respectfully requested.

If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5681-75600/RCK.

Respectfully submitted,

/Robert C. Kowert/
Robert C. Kowert, Reg. #39,255
Attorney for Applicant

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8850

Date: October 6, 2008